NETWORK INFRASTRUCTURE

Network Intrusion Detection/Prevention Systems
&
Content Scanning Appliances

Version 8, Release 1

Supplement of Network Infrastructure STIG, V8R1

24 March 2010

Developed by DISA for the DoD

This page is intentionally blank.

TABLE OF CONTENTS

Page

TABLE OF FIGURES

SUMMARY OF CHANGES

GENERAL CHANGES:

The previous STIG release was Version 7, Release 1, dated 25 October 2007.

This release concentrates on meeting the following objectives:

1) Sensor components and placement
2) Subnets requiring IDPS
3) Sensor Configuration
4) Protection from data overflow and scaling
5) Protecting sensor data in transit
6) Describing IP address assignment with illustration
7) Vulnerability ID renumbering structure

		V8 IDPS Detailed Changes		
STIG ID	V-Key	Short Name	Description of Change, Add or Delete in VMS	Network Security Checklist Addition or Deletion of vulnerability
NET-IDPS-001	V0018489	IDPS interface with data flow is not passive	New V8 Policy	IDS
NET-IDPS-002	V0018484	IDPS management servers located on MGT network	New V8 Policy	IDS
NET-IDPS-003	V0003179	IDPS data is not spooled before data overflow	New V8 Policy	IDS
NET-IDPS-004	V0018501	Messages do not indicate capacity is exceeded	New V8 Policy	IDS
NET-IDPS-005	V0018502	Whitelists and blacklists are not validated	New V8 Policy	IDS
NET-IDPS-006	V0018508	Web IDPS is not configured to protect web servers	New V8 Policy	IDS
NET-IDPS-007	V0018509	File Servers are not protected by IDPS	New V8 Policy	IDS
NET-IDPS-008	V0018513	IDPS does not protect IP hijacking TCP sessions	New V8 Policy	IDS
NET-IDPS-009	V0018512	IDPS is not configured to protect MGT network	New V8 Policy	IDS
NET-IDPS-010	V0019233	IDPS does not drop half open TCP sessions	New V8 Policy	IDS
NET-IDPS-011	V0019246	IDPS does not implement forged tcp resets	New V8 Policy	IDS
NET-IDPS-012	V0019250	IDPS does not protect from LAND DoS attack	New V8 Policy	IDS

V8 IDPS Detailed Changes				
STIG ID	V-Key	Short Name	Description of Change, Add or Delete in VMS	Network Security Checklist Addition or Deletion of vulnerability
NET-IDPS-013	V0019256	IDPS Atomic Signatures are not protecting enclave	New V8 Policy	IDS
NET-IDPS-016	V0018490	IDPS sensor is not monitoring DMZ segments	New V8 Policy	Network Policy for IDS
NET-IDPS-017	V0018491	IDPS is not monitoring VPN termination points	New V8 Policy	Network Policy for IDS
NET-IDPS-018	V0018492	IDPS sensor is not monitoring Server Farm segments	New V8 Policy	Network Policy for IDS
NET-IDPS-019	V0018493	IDPS sensor is not monitoring Network MGT network	New V8 Policy	Network Policy for IDS
NET-IDPS-020	V0018494	IDPS sensor is not monitoring WAN junction points	New V8 Policy	Network Policy for IDS
NET-IDPS-021	V0008272	IDPS is not monitoring traffic at the perimeter	IDS-Renamed from NET1330. Removed verbiage from longname requiring monitoring of all connections	Network Policy for IDS
NET-IDPS-022	V0014732	IDS is not NIAP approved	IDS-Renamed from NET1331. Removed EAL language.	Network Policy for IDS
NET-IDPS-023	V0018495	Remote IDS data is not collected by enterprise	New V8 Policy	Network Policy for IDS
NET-IDPS-024	V0018496	IDS traffic in transit is transmitted unprotected	New V8 Policy	Network Policy for IDS
NET-IDPS-025	V0003179	IDPS data from agent to MGT network is not secured	New V8 Policy	Network Policy for IDS
NET-IDPS-026	V0018503	Thresholds are not reviewed regularly	New V8 Policy	Network Policy for IDS
NET-IDPS-027	V0018504	Anomaly baselines are not periodically rebuilt	New V8 Policy	Network Policy for IDS
NET-IDPS-028	V0018505	Local networks do not get regular IDS updates	New V8 Policy	Network Policy for IDS
NET-IDPS-029	V0018506	Server is not configured to allow read-only access	New V8 Policy	Network Policy for IDS

V8 IDPS Detailed Changes				
STIG ID	V-Key	Short Name	Description of Change, Add or Delete in VMS	Network Security Checklist Addition or Deletion of vulnerability
NET-IDPS-030	V0018507	Access to signatures is not restricted	New V8 Policy	Network Policy for IDS
NET-IDPS-031	V0018510	Backups are not taken before updates	New V8 Policy	Network Policy for IDS
NET-IDPS-032	V0018511	Update files are not validated for accuracy	New V8 Policy	Network Policy for IDS
NET-IDPS-033	V0008078	NSO has not established weekly backup procedures	Renamed from net1346	Network Policy for IDS
NET-IDPS-034	V0008079	IDS Anti-virus updates procedures not in SOP	Renamed from net1348	Network Policy for IDS
NET-IDPS-035	V0008080	SA has not subscribed to the vendor notifications.	Renamed from net1350	Network Policy for IDS
NET-IDPS-036	V0015424	IDS software is not kept current	Renamed from net1351	Network Policy for IDS
NET-SRVFRM-001	V0019187	Servers do not employ HIDs.	New V8 Policy	Network Infrastructure Policy

1. NETWORK INTRUSION DETECTION

1.1 Components of Network Intrusion Detection/Prevention

An intrusion detection system (IDS) is software that automates the intrusion detection process. An intrusion prevention system (IPS) is software that has all the capabilities of an intrusion detection system and can also attempt to stop possible incidents. IDS and IPS technologies offer many of the same capabilities, and administrators can usually disable prevention features in IPS products, causing them to function as IDSs. Accordingly, for brevity the term intrusion detection and prevention system (IDPS) is used throughout this guide to refer to both IDS and IPS technologies.

All DoD locations will install, maintain, and operate an IDPS inside of their network enclaves. The Enclave IDPS will monitor internal network traffic and provide near real-time alarms for network-based attacks. A host intrusion detection (HID) application is not required on an OS-based IDPS.

Sensors and agents monitor and analyze activity. The term *sensor* is typically used for IDPSs that monitor networks, including network-based, wireless, and network behavior analysis technologies. The term *agent* is typically used for host-based IDPS technologies.

A *management server* is a centralized device that receives information from the sensors or agents and manages them. Some management servers perform analysis on the event information that the sensors or agents provide and can identify events that the individual sensors or agents cannot. Matching event information from multiple sensors or agents, such as finding events triggered by the same IP address, is known as *correlation*. Management servers are available as both appliance and software-only products. Some small IDPS deployments do not use any management servers, but most IDPS deployments do. In larger IDPS deployments, there are often multiple management servers, and in some cases there are two tiers of management servers.

Securing IDPS components is very important because IDPSs are often targeted by attackers. If an attacker can compromise an IDPS, it can be rendered useless in detecting subsequent attacks against other hosts. Also, IDPSs often contain sensitive information such as host configurations and known vulnerabilities that could be helpful in planning additional attacks. In addition to hardening software-based IDPS components and ensuring that all IDPS components are fully up-to-date, administrators should perform additional actions to ensure that the IDPS components themselves are secured appropriately. Administrators should create separate accounts for each user and administrator of the IDPS, and assign each account only the necessary privileges. Administrators should ensure that all IDPS management communications are protected appropriately, either through physical (e.g., management network) or logical (e.g., management VLAN) separation, or through encryption of communications. If encryption is used for protection, it should be performed using FIPS-approved encryption algorithms.

The site may establish a support agreement with the Computer Network Defense Service Provider (CNDSP) for monitoring. The local staff is responsible for initial response to real-time alarms. A Memorandum of Understanding (MOU) should be in place that documents the support agreement.

1.2 IP Assignments on Network Sensors

Administrators should ensure that for both passive and inline sensors, IP addresses are not assigned to the network interfaces used to monitor network traffic, except for network interfaces used for IDPS management. Operating a sensor without IP addresses assigned to its monitoring interfaces is known as operating in stealth mode. Stealth mode improves the security of the IDPS sensors because it prevents other hosts from initiating connections to them. This conceals the sensors from attackers and thus limits their exposure to attacks. If monitoring is being performed using a switch SPAN port, it is recommended that the IDPS is configured in Stealth Mode; the network interface card (NIC) connected to the SPAN port would not have any network protocol stacks bound to it. A second NIC would then be connected to an Out of Band (OOB) network. Stealth mode will reduce the risk of the IDPS itself being attacked.

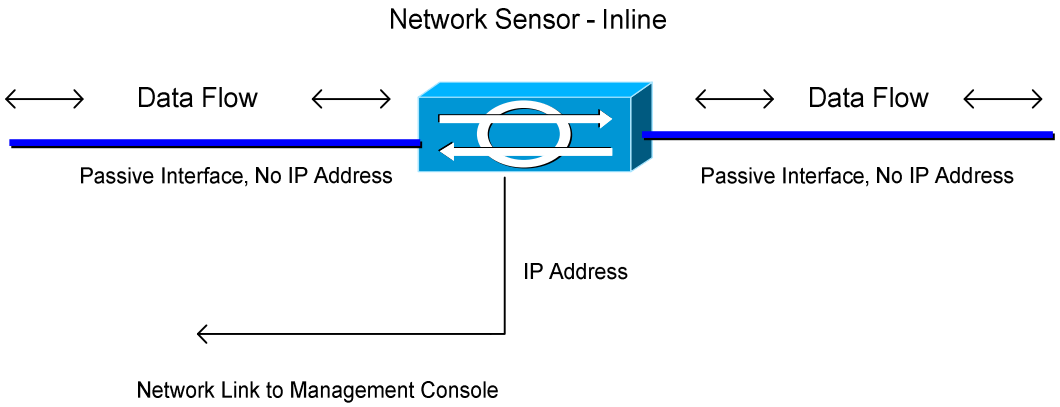

Figure 1-1. Network Sensor - Inline

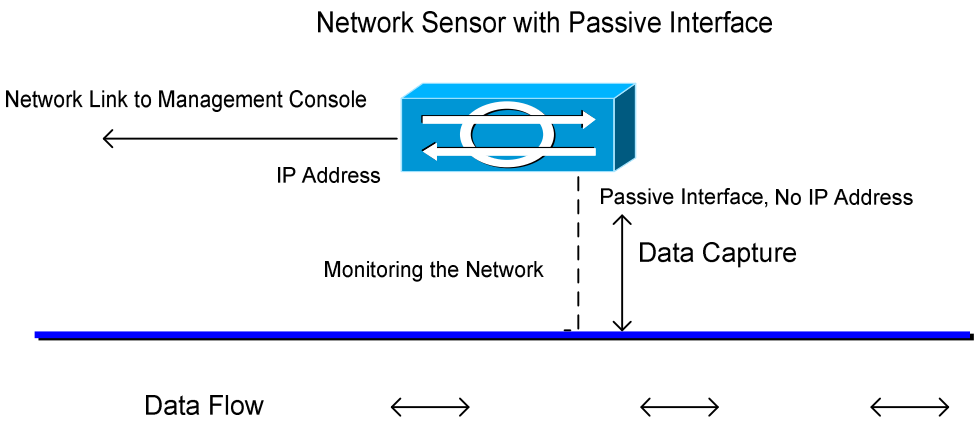

Figure 1-2. Network Sensor with Passive Interface

1.3 Network Sensor

The initial step in IDPS deployment is determining where sensors should be placed. Because attacks originate at the enclave perimeter and within the enclave boundary an IDPS implementation at the enclave perimeter only will not suffice. By placing IDPS technology throughout the Enterprise Regional enclaves and stand-alone enclaves, system administrators can track the spread of attacks and take corrective actions to prevent attacks reaching critical resources. The following locations should have sensor deployments in the Regional enclave:

1. An IPS installed to protect databases from known network and SQL specific attacks to prevent behaviors such as a buffer overflow, worms and other attacks targeting known vulnerabilities in database platforms. Black and whitelisting of protocols and IP address combinations can also be applied.

2. IDS behind the perimeter firewall

3. On demilitarized zone (DMZ) segments that house public servers (Web, Secure File Transfer Protocol [SFTP], Domain Name System [DNS]), email gateways, etc)

4. Behind VPN concentrators to monitor unencrypted VPN traffic and behind all tunnel endpoints to monitor all traffic (IPv4 and IPv6) entering the enclave

5. On segments that house intranet services that are sensitive according to the defined security policy and on critical resource segments (Server Farms segments containing databases, private backend servers, personnel data, etc)

6. On segments that house network and security management servers (Network Management segments or OOB networks)

7. At WAN junction points between the Regional enclave and the local enclave networks as well as between the enterprise Regional enclave and tenant network enclaves

The following locations should have sensor deployments in the Local enclave (base, camp, post, and station) at a minimum:

1. IDS behind the perimeter firewall

If the Local enclave provides services typically found at the regional enclave (tunnels, DMZs, Server Farms, etc), sensors will be deployed to monitor those traffic endpoints.

The server farm is often overlooked from a security perspective. When examining the levels of access most employees have to the servers to which they attach, the servers can often become the primary goal of internally originated attacks. Simply relying on effective passwords does not provide for a comprehensive attack mitigation strategy. Using host and network-based IDS, private VLANs, access control, and good system administration practices (such as keeping systems up to date with the latest patches), provides a much more comprehensive response to attacks.

1.4 Regional Enclave - Sensor Data logs in Transit

The enterprise Regional Enclave will develop a hierarchical monitoring structure that allows the captured local enclave (base, camp, post, and station) traffic to be exported to the regional enclave for trend analysis and reporting. The local enclave sensor data should be protected at all time in transit to the Regional Enclave (NOC / Data Center) by an OOB network or authenticated tunnel. All IDPS data collected by agents in the enclave at required locations must also be protected by logical separation when in transit from the agent to the management or database servers located on the Network Management subnet.

1.5 Scaling Sensors – Data Overflow Protection

Events on the sensor are typically stored on a large input queue. The queue in the sensor is typically very large and can hold several days of logging events under normal conditions. Nevertheless, the monitoring application must retrieve events from the sensor before the queue becomes full; otherwise the sensor will start overwriting the unread events.

Scaling IDPS sensors to avoid missed packets as a result of CPU and memory thresholds when link mbps is greater than what the engine can inspect should be an initial consideration prior to deployment. The IDPS administrator will have the sensor send notifications to the syslog server or central controller when thresholds limits do occur.

1.6 Configuration Policies for Sensors

The IDS or firewall is the first device that is under the sites control that has the possibility to alarm the local staff of an ongoing attack. An alert from either of these devices can be the first indication of an attack or system failure.

A blacklist is a list of discrete entities, such as hosts, TCP or UDP port numbers, ICMP types and codes, applications, usernames, URLs, filenames, or file extensions, that have been previously determined to be associated with malicious activity. Blacklists, also known as hot lists, are typically used to allow IDPSs to recognize and block activity that is highly likely to be malicious, and may also be used to assign a higher priority to alerts that match entries on the blacklists. Some IDPSs generate dynamic blacklists that are used to temporarily block recently detected threats (e.g., activity from an attacker's IP address). A whitelist is a list of discrete entities that are known to be benign. Whitelists are typically used on a granular basis, such as protocol-by-protocol, to reduce or ignore false positives involving known benign activity from trusted hosts. Whitelists and blacklists are most commonly used in signature-based detection and stateful protocol analysis.

Administrators should review tuning and customizations periodically to ensure that they are still accurate. For example, whitelists and blacklists should be checked regularly and all entries validated to ensure that they are still accurate and necessary. Thresholds and alert settings might need to be adjusted periodically to compensate for changes in the environment and in threats. Edits to detection code might need to be replicated whenever the product is updated (e.g., patched, upgraded). Administrators should also ensure that any products collecting baselines for anomaly-based detection have their baselines rebuilt periodically as needed to support accurate detection. The IAM is required to have the enclave prepared for readiness by raising INFOCON levels prior to an activity to ensure the network is as ready as possible when the operation or exercise begins. Because system and network administrators implement many of the INFOCON measures over a period of time in a pre-determined operational rhythm, commanders should raise INFOCON levels early enough to ensure completion of at least one cycle before the operational activity begins. Recommendations for possible INFOCON changes should be written into Operation Plans (OPLAN) and Concept Plans (CONPLAN). Guidelines can be found in Strategic Command Directive (SD) 527-1.

In a large scale IDPS deployment, it is common to have an automated update process implemented. This is accomplished by having the updates downloaded on a dedicated SFTP server within the management network. The SFTP server should be configured to allow read-only access to the files within the directory on which the signature packs are placed, and then only from the account that the sensors will use. The sensors can then be configured to automatically check the SFTP server periodically to look for the new signature packs and to update themselves once they have been tested.

1.7 Content Scanning Policies

In the Regional Enterprise Enclave different sets of sensors will see different traffic as a result of their location within the regional enclave. By establishing separate signature profiles for each set of sensors, each profile can then be tuned to generate alarms based on the traffic types seen, the attack signatures, and the specific traffic (string signatures) that is relevant to that particular set of sensors. If more than one set of sensors will see the same traffic types, then the same signature profile may be used for both sets. Alerting on specific connection signatures, general attack signatures, and specific string signatures provides focused segment analysis at Layers 4 through 7.

The IDPS system administrator will ensure the sensor monitoring the web servers is configured for application inspection and control of all web ports e.g. 80, 3128, 8000, 8010, 8080, 8888, 24326, etc. The sensor monitoring the web servers should be capable of inspecting web traffic that is not received on web ports; known as port redirection. In many implementations this is a separate signature that needs to be enabled.

Network segments containing FTP servers should have sensors installed that monitor, inspect and log all recognized ftp commands, as well as unrecognized ftp commands.

There are a number of publicly available tools that exist to facilitate the hijacking of TCP sessions. An attacker using such tools can determine the TCP sequence and acknowledgement numbers that two hosts are using in a communication session. This information could enable the attacker to take over the legitimate network connection of an authorized user and inject commands into the session. This is particularly serious because most forms of one-time passwords do not prevent this access.

The management network must detect all attacks on the management hosts. The management network has a various range of traffic that is permitted. Some of the following traffic is allowed on the Management Hosts Segment: *Trivial File Transfer Protocol (TFTP [UDP 69])*—For network device configuration files from devices on the Managed Devices Segment; *FTP-Data (TCP 20)*—For file transfers to network devices on the Managed Devices Segment and for Internet downloads; *FTP-Control (TCP 21)*—For file transfers to network devices on the Managed Devices Segment and for Internet downloads; *Syslog (UDP 514)*—From network devices on the Managed Devices Segment; *Telnet (TCP 23)*—To network devices on the Managed Devices Segment; *SSH (TCP 22)*—To network devices on the Managed Devices Segment; *Network Time Protocol (NTP [UDP 123])*—To synchronize the clocks of all network devices on the Managed Devices Segment; *HTTP (TCP 80)*—To the Internet and from hosts on other segments to download the host-based IPS agent software; *HTTPS (TCP 443)*—To network devices on the Managed Devices Segment and the Internet as well as between the host-based IPS Console and its agents; *TACACS+ (TCP 49)*—For administrator authentication to devices on the Managed Devices Segment; *RADIUS (UDP 1812/1813 authentication/accounting)*—For authentication of administrator remote-access VPN connections coming from the Remote Administration Segment; *ICMP (IP Protocol 1)*—Echo request and response to reach network devices on the Managed Devices Segment and the Internet; *DNS (UDP 53)*—For name

translation services for management hosts as they access services on the Internet; *Simple Network Management Protocol (SNMP [UDP 161])*—To query information from network devices on the Managed Devices Segment; *SNMP-Trap (UDP 162)*—To receive trap information from network devices on the Managed Devices Segment.

SYN flood sends a flood of TCP/SYN packets, often with a forged sender address. Each of these packets is handled like a connection request, causing the server to spawn a half-open connection, by sending back a TCP/SYN-ACK packet, and waiting for a packet in response from the sender address. However, because the sender address is forged, the response never comes. These half open connections saturate the number of available connections the server is able to make, keeping it from responding to legitimate requests until after the attack ends.

By listening to the conversation flow of inbound and outbound internet traffic for malware and malware references, the IDPS can prevent unwanted programs entering into the enclave. When it detects unmanaged instant messaging and peer-to-peer protocols or malware coming over IM, the IDPS can prevent the unwanted computer programs from entering the network by spoofing the source and destination machine addresses to send each session partner a TCP Reset packet. The TCP Reset instructs both sender and receiver to cease the current transfer of data.

The LAND attack is a denial-of-service (DoS) attack in which an attacker sends a TCP packet (with the SYN bit set) to a system in which the source and destination IP address (along with the source and destination port) are the same. When it was first discovered, many IP stacks crashed the system when they received a LAND attack. An effective implementation is the use of an Atomic attack signature that looks at a single packet, because State information (tracking established connections) is not necessary in identifying this attack.

Without an industry agreed-upon set of definitions for IDPS controls, the use of the term signature will apply to all IDPS technologies. Signatures are defined as identifying something, defining it and then stop it from occurring. Signatures fall into one of the following two basic categories depending on their functionality: Atomic or Stateful. Atomic signatures trigger on a single event, they do not require your intrusion system to maintain state. The entire inspection can be accomplished in an atomic operation that does not require any knowledge of past or future activities. These signatures consume minimal resources (such as memory) on the IPS/IDS device. These signatures are easy to understand because they search only for a specific event.

1.8 Software Updates and Signature Updates

There are two types of IDPS updates: software updates and signature updates. Software updates fix bugs in the IDPS software or add new functionality, while signature updates add new detection capabilities or refine existing detection capabilities (e.g., reducing false positives). For many IDPSs, signature updates cause program code to be altered or replaced, so they are really a specialized form of software update. For other IDPSs, signatures are not written in code, so a signature update is a change to the configuration data for the IDPS.

Software updates can include any or all IDPS components, including sensors, agents, management servers, and consoles. Software updates for sensors and management servers,

particularly appliance-based devices, are often applied by replacing an existing IDPS CD with a new one and rebooting the device. Many IDPSs run the software directly from the CD, so that no software installation is required. Other components, such as agents, require an administrator to install software or apply patches, either manually on each host or automatically through IDPS management software. Some vendors make software and signature updates available for download from their Web sites or other servers; often, the administrator interfaces for IDPSs have features for downloading and installing such updates.

Administrators should verify the integrity of updates before applying them, because updates could have been inadvertently or intentionally altered or replaced. The recommended verification method depends on the update's format, as follows:

Files downloaded from a Web site or FTP site. Administrators should compare file checksums provided by the vendor with checksums that they compute for the downloaded files.

Update downloaded automatically through the IDPS user interface. If an update is downloaded as a single file or a set of files, either checksums provided by the vendor should be compared to checksums generated by the administrator, or the IDPS user interface itself should perform some sort of integrity check. In some cases, updates might be downloaded and installed as one action, precluding checksum verification; the IDPS user interface should check each update's integrity as part of this.

Removable media (e.g., CD, DVD). Vendors may not provide a specific method for customers to verify the legitimacy of removable media apparently sent by the vendors. If media verification is a concern, administrators should contact their vendors to determine how the media can be verified, such as comparing vendor-provided checksums to checksums computed for files on the media, or verifying digital signatures on the media's contents to ensure they are valid. Administrators should also consider scanning the media for malware, with the caveat that false positives might be triggered by IDPS signatures for malware on the media.

2. EXTERNAL INTRUSION DETECTION

A CND device should be positioned to best detect any traffic crossing the boundary between the enclave and the GiG or non-DoD source. CND devices will be configured in accordance with an accredited Tier II CNDSP's guidance and should be configured as a minimum to pass reported unauthorized or suspicious traffic to an accredited Tier II CNDSP. Reference the diagram in the AG section for AG connectivity.

Placing the CND device on the exterior between the premise router and the NIPRNet / SIPRNet customer edge (CE) node router will enable the CNDSP to detect attempted attacks that may otherwise be blocked by the premise router or firewall.

A signature-based, anomaly-based, or rules-based CND device that has been customized to specific NIPRNet or SIPRNet traffic can alert CND Service Provider of suspected threats at the enclave's gateway. Enclave security controls will be configured to allow the Tier II CNDSP to position, configure, and maintain their view of any CND devices in accordance with a mutually agreed upon MOA.

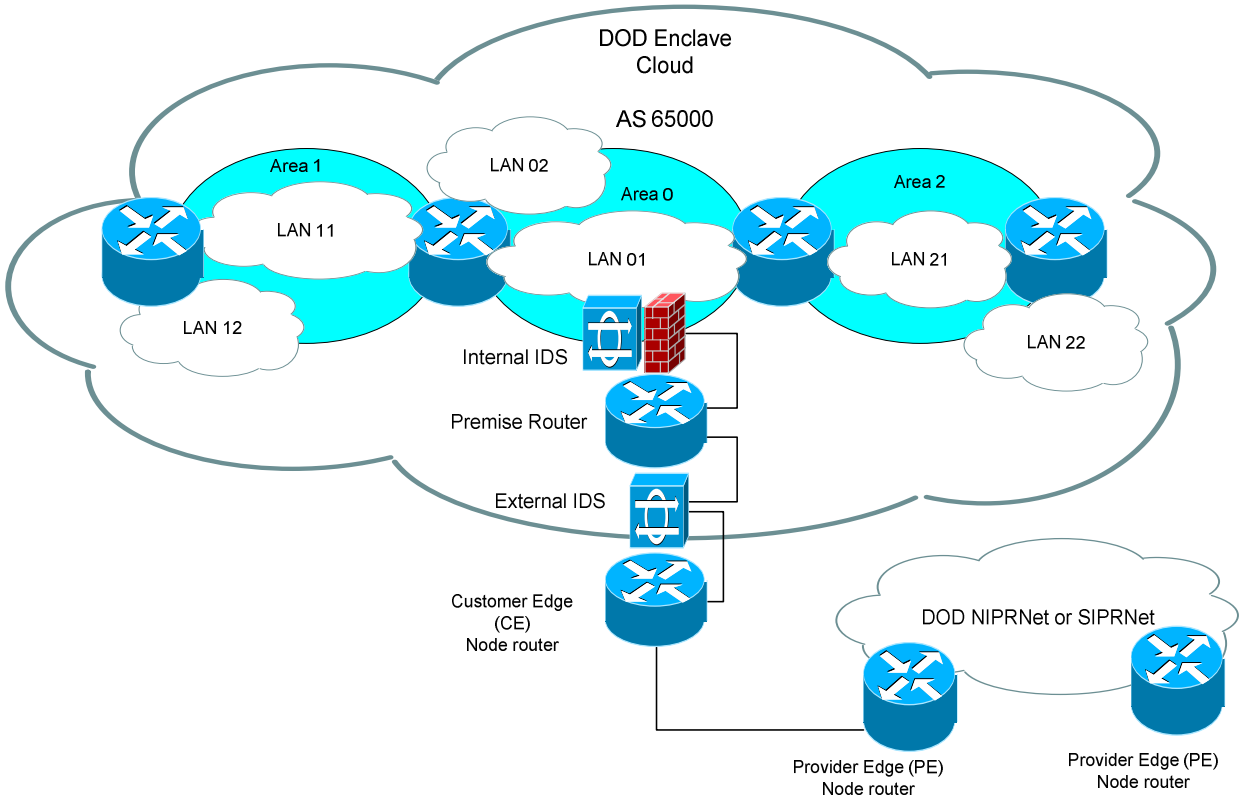

Figure 2-1. External IDS

This page is intentionally blank.